

Simultaneous sparse estimation of canonical vectors in the $p \gg N$ setting

Irina Gaynanova*, James G. Booth[†] and Martin T. Wells*

Abstract

This work considers the problem of sparse estimation of canonical vectors in linear discriminant analysis when $p \gg N$. Several methods have been proposed in the literature that estimate one canonical vector in the two-group case. However, $G - 1$ canonical vectors can be considered if the number of groups is G . In the multi-group context, it is common to estimate canonical vectors in a sequential fashion. Moreover, separate prior estimation of the covariance structure is often required. In contrast, the proposed method estimates all canonical vectors directly. First, we highlight the fact that in the $N > p$ setting the canonical vectors can be expressed in a closed form up to an orthogonal transformation. Secondly, we propose an extension of this form to the $p \gg N$ setting and achieve feature selection by using a group penalty. The resulting optimization problem is convex and can be solved using block-coordinate descent algorithm. Furthermore, we show that the proposed method can consistently identify the true support of canonical vectors. The practical performance of the method is evaluated through simulation studies as well as real data applications.

Keywords: Block-coordinate descent; Classification; Dimension reduction; Discriminant analysis; Feature selection; Group penalization.

1 Introduction

Recent technological advances have generated high-dimensional data sets across a wide variety of application areas including finance, atmospheric science, astronomy, biology and medicine. Not only do these data sets provide computational challenges, but they also pose new statistical challenges as the standard methods no longer apply. Linear Discriminant Analysis (LDA) is one of the popular classification and data visualization tools that is used in the $N \gg p$ setting. LDA seeks the linear combinations of features that maximize between group variability with respect to within group variability (Mardia et al., 1979, Chapter 11). These linear combinations are called *canonical vectors* and they provide a low-dimensional

*Department of Statistical Science, Cornell University, Ithaca, NY; e-mail: ig93@cornell.edu

[†]Department of Biological Statistics and Computational Biology, Cornell University, Ithaca, NY

This research was partially supported by NSF-DMS 1208488 and ASAF grant FA9550-13-1-0137. We are grateful to Jacob Bien for a valuable discussion of block-coordinate descent algorithms.

representation of the data by reducing the original feature space dimension p to $G - 1$, where G is the total number of groups.

The classical use of LDA in the setting when $p \gg N$ fails to provide good results, largely due to the singularity of covariance matrix and over-selection of relevant features (Dudoit et al., 2002; Bickel & Levina, 2004). As a result, the extension of LDA to high-dimensional settings has recently received a lot of attention in the literature. A number of these proposals give rise to non-sparse classifiers. Friedman (1989), Krzanowski et al. (1995) and Xu et al. (2009) regularize the within-class covariance matrix in order to obtain a positive definite estimate. Other approaches that lead to sparse discriminant vectors have also been considered. Tibshirani et al. (2002) propose the shrunken centroids methodology by adapting the naive Bayes classifier and soft-thresholding the mean vectors. Guo et al. (2007) combine the shrunken centroids approach with a ridge-type penalty on the within-class covariance matrix. Witten & Tibshirani (2011) apply an ℓ_1 penalty to the Fisher’s discriminant problem in order to obtain sparse discriminant vectors. Clemmensen et al. (2011) use an optimal scoring approach which essentially reduces the sparse discriminant vector construction to a penalized regression problem.

In the two-group setting, Cai & Liu (2011) and Mai et al. (2012) propose direct estimation of the canonical vector thus avoiding separate estimation of the covariance matrix. Simulations and real data applications show that the direct estimation approach results in improved misclassification rates in comparison to alternative methods. Moreover, the corresponding optimization problems can be solved efficiently for large data sets and have desirable theoretical properties. Unfortunately, the extension of the two-group methods to the multi-group case is ambiguous (Hastie et al., 2009, p. 658). Popular approaches include “one-versus-all” and “one-versus-one” methods, where the final classification assignment is usually based on the “majority vote”. As such, computation of a larger number of vectors is required (G and $G(G - 1)/2$ versus $G - 1$).

Witten & Tibshirani (2011) and Clemmensen et al. (2011) propose estimating canonical vectors in a sequential fashion in the multi-group setting: starting with the first canonical vector v_1 , with subsequent v_i found subject to orthogonality constraints. This approach is undesirable not only from the computational viewpoint, but also from the estimation perspective. Each subsequent canonical vector v_i relies on all the previous estimates v_k for $k < i$, hence there is a propagation of the estimation error. Moreover, the corresponding optimization problems are nonconvex which means that the convergence of the optimization algorithms to the global solution is not guaranteed. This computational burden poses additional theoretical challenges in the analysis and, as a result, these methods do not come with theoretical guarantees.

The objective of this article is to bridge the computational and theoretical gap in the literature between the two-group and multi-group methods. Inspired by the superior performance of the proposals by Cai & Liu (2011) and Mai et al. (2012) in the two-group case, our goal is to introduce and develop a novel methodology that has such guaranteed performance in the multi-group setting. Our proposal is based on the observation that canonical vectors can be expressed in a closed form up to an orthogonal transformation. Moreover, this transformation affects neither the classification rule nor the sparsity pattern. As a result, we propose an algorithm that estimates all $G - 1$ canonical vectors at once without prior estimation of the covariance structure. To our knowledge, this is the first method that yields

both variable selection and direct estimation of canonical vectors in the multi-group setting. The proposed optimization problem is convex and therefore can be solved efficiently for large data sets. The algorithm doesn't require additional regularization of the sample covariance matrix or generation of an initial starting point. Furthermore, the form of the optimization problem suggests a natural choice for the grid of tuning parameters. This is not the case for the methods of [Cai & Liu \(2011\)](#), [Witten & Tibshirani \(2011\)](#) and [Clemmensen et al. \(2011\)](#) where the appropriate grid needs to be carefully chosen by the user. While these advantages are mostly computational, they simplify the implementation of the method leading to more consistent results across users. More implementation details are provided in [Section 4.3](#).

Though the motivation for our approach is quite different from [Mai et al. \(2012\)](#), we show that the two methods are equivalent in the two-group setting. This connection allows us to extend the theoretical results of [Mai et al. \(2012\)](#) to the multi-group setting and hence to show that the proposed method can consistently identify the true support of canonical vectors.

The rest of the paper is organized as follows. In [Section 2](#) we discuss the canonical vectors estimation problem and provide new insights into the form of the canonical vectors in the standard $N \gg p$ setting. We use these insights to propose a direct estimation procedure and describe computational aspects of the algorithm. In [Section 3](#) we develop bounds on the estimation error and prove that the proposed method identifies the true support of canonical vectors with high probability. [Section 4](#) provides simulation results while [Section 5](#) describes applications to real datasets. We conclude with some discussion in [Section 6](#).

2 Methodology

2.1 Notation

For a vector $b \in \mathbb{R}^p$ we define $\|b\|_\infty = \max_{i=1,\dots,p} |b_i|$, $\|b\|_1 = \sum_{i=1}^p |b_i|$ and $\|b\|_2 = \sqrt{\sum_{i=1}^p b_i^2}$. For a matrix $M \in \mathbb{R}^{n \times p}$ we define m_i to be its i th row and M_j to be its j th column. We also define $\|M\|_\infty = \max_{i=1,\dots,n} \|m_i\|_1$, $\|M\|_{\infty,2} = \max_{i=1,\dots,n} \|m_i\|_2$, $\|M\|_1 = \sum_{i=1}^n \sum_{j=1}^p |m_{ij}|$, $\|M\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^p m_{ij}^2}$ and $\|M\|_* = \sum_{i=1}^{\min(n,p)} \sigma_i(M)$, where $\sigma_i(M)$ is the i th singular value of M . We define \mathbb{O}^p to be the space of $p \times p$ orthogonal matrices R such that $RR^t = R^tR = I$.

2.2 Estimation problem

We assume that $X_i \in \mathbb{R}^p$, $i = 1, \dots, N$, are independent and come from G groups with different means and the same covariance matrix, i.e. $\mathbb{E}(X_i|Y_i = g) = \mu_g$ and $\text{Cov}(X_i|Y_i = g) = \Sigma_W$, where $Y_i \in \{1, \dots, G\}$. The between-group population covariance matrix Σ_B is defined as

$$\Sigma_B = \sum_{g=1}^G \pi_g (\mu_g - \mu)(\mu_g - \mu)^t,$$

where $\pi_g = P(Y_i = g)$ are group-specific probabilities and $\mu = \sum_{i=1}^G \pi_g \mu_g$ is the overall population mean. The *population canonical vectors* Ψ are defined as eigenvectors corresponding

to non-zero eigenvalues of $\Sigma_W^{-1}\Sigma_B$ (unique up to a normalization). For a new observation value of X , $x \in \mathbb{R}^p$, the population classification rule $h_\Psi(x)$ is defined as

$$h_\Psi(x) = \arg \min_{1 \leq g \leq G} (x - \mu_g)^t \Psi (\Psi^t \Sigma_W \Psi)^{-1} \Psi^t (x - \mu_g) - 2 \log \pi_g. \quad (1)$$

Our goal is to estimate Ψ based on the sample observations $X_i \in \mathbb{R}^p$ and sample labels $Y_i \in \{1, \dots, G\}$.

Consider the within-group sample covariance matrix $W = \frac{1}{N-G} \sum_{g=1}^G (n_g - 1) S_g$ and the between-group sample covariance matrix $B = \frac{1}{N} \sum_{g=1}^G n_g (\bar{X}_g - \bar{X})(\bar{X}_g - \bar{X})^t$, where n_g is the number of observations in group g , S_g is the sample covariance matrix for group g , \bar{X}_g is the sample mean for group g and \bar{X} is the overall sample mean. Recall that W is nonsingular when $N \gg p$ and therefore we can define the *sample canonical vectors* V as $G - 1$ eigenvectors corresponding to non-zero eigenvalues of $W^{-1}B$ (Mardia et al., 1979, Chapter 11.5). Similarly to (1), the sample classification rule $\hat{h}_V(x)$ is defined as

$$\hat{h}_V(x) = \arg \min_{1 \leq g \leq G} (x - \bar{X}_g)^t V (V^t W V)^{-1} V^t (x - \bar{X}_g) - 2 \log \frac{n_g}{N}. \quad (2)$$

In what follows we show that in the $N \gg p$ setting, canonical vectors can be expressed in a closed form up to an orthogonal transformation (Proposition 2). Since this transformation has no effect on the classification rule (Proposition 3), it allows us to estimate Ψ directly.

Proposition 1. *The following decompositions hold: $\Sigma_B = \Delta \Delta^t$ and $B = D D^t$, where for $r = 1, \dots, G - 1$ the r th column of Δ has the form*

$$\Delta_r = \frac{\sqrt{\pi_{r+1}} (\sum_{i=1}^r \pi_i (\mu_i - \mu_{r+1}))}{\sqrt{\sum_{i=1}^r \pi_i \sum_{i=1}^{r+1} \pi_i}} \quad (3)$$

and the r th column of D has the form

$$D_r = \frac{\sqrt{n_{r+1}} (\sum_{i=1}^r n_i (\bar{X}_i - \bar{X}_{r+1}))}{\sqrt{N} \sqrt{\sum_{i=1}^r n_i \sum_{i=1}^{r+1} n_i}}. \quad (4)$$

Notice that the columns specify orthogonal contrasts between the means of G groups. Also, in the case $G = 2$, $\Delta = \sqrt{\pi_1 \pi_2} (\mu_2 - \mu_1)$ and $D = \frac{\sqrt{n_1 n_2}}{N} (\bar{X}_2 - \bar{X}_1)$.

Proposition 2. *Define Δ and D as in (3) and (4). There exists a matrix $P \in \mathbb{O}^{G-1}$ such that $\Psi = \Sigma_W^{-1} \Delta P$. Moreover, if W is nonsingular, there exists a matrix $R \in \mathbb{O}^{G-1}$ such that $V = W^{-1} D R$.*

Proposition 3. *The population classification rule based on Ψ is the same as the population classification rule based on $\tilde{\Psi} = \Sigma_W^{-1} \Delta$: $h_\Psi(x) = h_{\tilde{\Psi}}(x)$ for all $x \in \mathbb{R}^p$. If W is nonsingular, then $\hat{h}_V(x) = \hat{h}_{W^{-1}D}(x)$ for all $x \in \mathbb{R}^p$.*

2.3 Proposed Estimation Criterion

From Propositions 2 and 3 it follows that for classification it is sufficient to estimate

$$\tilde{\Psi} = \Sigma_W^{-1} \Delta \quad (5)$$

rather than Ψ . To illustrate the motivation behind the proposed optimization problem, we first discuss our choice of the loss function and then our choice of the penalty.

Our first goal is to choose a suitable loss function that will capture the deviations of the estimator from the target $\tilde{\Psi} = \Sigma_W^{-1} \Delta$. We note that $\tilde{\Psi}$ in (5) can be defined as

$$\tilde{\Psi} = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \|\Sigma_W^{1/2} V - \Sigma_W^{-1/2} \Delta\|_F^2 = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \text{Tr} (V^t \Sigma_W V - 2 \Delta^t V). \quad (6)$$

In the two-group case, V is a vector and the objective function in (6) reduces to

$$\frac{1}{2} (\Sigma_W^{1/2} V - \Sigma_W^{-1/2} \Delta)^t (\Sigma_W^{1/2} V - \Sigma_W^{-1/2} \Delta) = \frac{1}{2} (V - \Sigma_W^{-1} \Delta)^t \Sigma_W (V - \Sigma_W^{-1} \Delta) = \frac{1}{2} (V - \tilde{\Psi})^t \Sigma_W (V - \tilde{\Psi}).$$

This objective function is the same as the quadratic loss function considered by Rukhin (1992), who observed that it is invariant with respect to linear transformation of the data. Hence, we can define an estimator \tilde{V} by substituting Σ_W and Δ with W and D :

$$\tilde{V} = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \text{Tr} (V^t W V - 2 D^t V). \quad (7)$$

Our second goal is to perform a variable selection, which in the discriminant analysis framework corresponds to estimating some entries of $\tilde{\Psi}$ exactly as zero. A common penalty that is used in this context is an ℓ_1 penalty $\|V\|_1 = \sum_{i=1}^p \sum_{j=1}^{G-1} |v_{ij}|$. Although this penalty leads to sparse estimates, there is no guarantee that the features are simultaneously eliminated from all canonical vectors. In other words, although each canonical vector is sparse individually, the total number of features used may be very large. Moreover, the sparsity of the individual matrix elements is not preserved under the orthogonal rotation. Hence, the sparsity of elements of $\tilde{\Psi}$ doesn't imply the sparsity of elements of Ψ .

To overcome the overfitting and nonorthogonality issues, we consider the row-wise penalty $\sum_{i=1}^p \|v_i\|_2$. This penalty results in an estimate that is invariant to orthogonal transformation and eliminates features from all canonical vectors at once inducing row sparsity on the matrix V . Combining (7) with this penalty suggests an estimator $\hat{V}(\lambda)$, defined as

$$\hat{V}(\lambda) = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \text{Tr} (V^t W V - 2 D^t V) + \lambda \sum_{i=1}^p \|v_i\|_2. \quad (8)$$

When W is nonsingular and $\lambda = 0$, $\hat{V}(\lambda) = W^{-1} D$, which according to Proposition 2 is the matrix of sample canonical vectors up to an orthogonal rotation. Unfortunately, the objective function in (8) can be unbounded when W is singular since $\text{Tr} (V^t W V - 2 D^t V)$ can be made arbitrarily small. Hence, an additional regularization of (8) is required.

A simple solution is to use $\tilde{W} = W + \rho I$ instead of W in (8). This type of regularization is quite common in the LDA context and is for example used by [Friedman \(1989\)](#), [Guo et al. \(2007\)](#) and [Cai & Liu \(2011\)](#). In our case it leads to

$$\begin{aligned}\hat{V}(\lambda, \rho) &= \arg \min_{V \in \mathbb{R}^{p \times G-1}} \text{Tr} \left(\frac{1}{2} V^t \tilde{W} V - D^t V \right) + \lambda \sum_{i=1}^p \|v_i\|_2 \\ &= \arg \min_{V \in \mathbb{R}^{p \times G-1}} \frac{1}{2} \text{Tr}(V^t W V) + \frac{\rho}{2} \|V - D\|_F^2 + \lambda \sum_{i=1}^p \|v_i\|_2.\end{aligned}$$

The second component of the objective function encourages $\hat{V}(\lambda, \rho)$ to be close to D , especially when ρ is large. In contrast, $\hat{V}(\lambda, \rho)$ should be close to $W^{-1}D$ according to Proposition 2. This discrepancy suggests that strong regularization of W may have a negative affect on classification performance.

Instead, we propose to substitute W in (8) with $T = W + DD^t = W + B$. Such substitution is possible because the matrices $W^{-1}B$ and $(W + B)^{-1}B$ have the same eigenvectors corresponding to non-zero eigenvalues. This is formalized in Proposition 4 below. The resulting estimator has the form

$$\hat{V}(\lambda) = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \text{Tr}(V^t W V) + \frac{1}{2} \|D^t V - I\|_F^2 + \lambda \sum_{i=1}^p \|v_i\|_2, \quad (9)$$

where the objective function is convex and bounded below by zero. The three components of the objective function in (9) attempt to minimize the within-group variability, control the level of the between-group variability and provide regularization by inducing sparsity respectively.

Proposition 4. *Let Ψ be the matrix of eigenvectors corresponding to non-zero eigenvalues of $\Sigma_W^{-1} \Sigma_B$ and Υ_ρ be the matrix of eigenvectors corresponding to non-zero eigenvalues of $(\Sigma_W + \rho \Sigma_B)^{-1} \rho \Sigma_B$ for some positive ρ . Then there exists a diagonal matrix K_ρ such that $\Psi = \Upsilon_\rho K_\rho$.*

2.4 Connection with other sparse discriminant analysis methods when $G = 2$

The motivation for our method is based on the eigenstructure of the discriminant analysis problem in the multi-group setting, however it has a direct connection with the two-group methods previously proposed in the literature.

When $G = 2$, V is a vector in \mathbb{R}^p and (9) takes the form

$$\hat{V}(\lambda) = \arg \min_{V \in \mathbb{R}^p} \frac{1}{2} V^t W V + \frac{1}{2} (D^t V - 1)^2 + \lambda \|V\|_1.$$

Proposition 5. *Consider $\hat{V}_{DSDA}(\lambda)$ ([Mai et al., 2012](#)), defined as*

$$\hat{V}_{DSDA}(\lambda) = \arg \min_{\beta_0 \in \mathbb{R}, V \in \mathbb{R}^p} \frac{1}{2N} \sum_{i=1}^N (y_i - \beta_0 - X_i^t V)^2 + \lambda \|V\|_1,$$

where $y_i = -\frac{N}{n_1}$ if the i th subject is in group 1 and $y_i = \frac{N}{n_2}$ otherwise. Then

$$\hat{V}(\lambda) = \frac{N}{\sqrt{n_1 n_2}} \hat{V}_{DSDA} \left(\frac{N}{\sqrt{n_1 n_2}} \lambda \right).$$

Furthermore, [Mai & Zou \(2013\)](#) show an equivalence between the three methods for sparse discriminant analysis in the two-group setting: [Wu et al. \(2009\)](#), [Clemmensen et al. \(2011\)](#) and [Mai et al. \(2012\)](#). It follows that our method belongs to the same class, however it can be applied to any number of groups. Thus, it can be viewed as a multi-group generalization of this class of methods.

The optimization problem in (9) corresponds to the choice of $\rho = 1$ in Proposition 4. In general, any $\rho > 0$ leads to

$$\hat{V}(\lambda, \rho) = \arg \min_{V \in \mathbb{R}^{p \times G-1}} \frac{1}{2} \text{Tr}(V^t W V) + \frac{\rho}{2} \|D^t V - I\|_F^2 + \lambda \sum_{i=1}^p \|v_i\|_2. \quad (10)$$

When $\rho \rightarrow \infty$, (10) is equivalent to

$$\hat{V}(\lambda, \rho = \infty) = \arg \min_{D^t V = I} \frac{1}{2} \text{Tr}(V^t W V) + \lambda \sum_{i=1}^p \|v_i\|_2, \quad (11)$$

and therefore for large values of ρ the optimization problem (10) can be considered a convex relaxation to (11). When the number of groups is two, the optimization problem (10) is equivalent to the proposal of [Fan et al. \(2012\)](#), who also observe the connection between (10) and (11). They perform a simulation study to assess the effect of the tuning parameter ρ and note that its value doesn't significantly affect the classification results as long as the best λ is chosen for each ρ . They keep the value of ρ at a fixed level $\rho = 10$.

2.5 Optimization Algorithm

The optimization problem in (9) is convex with respect to V and therefore can be solved efficiently using, for example, a block-coordinate descent algorithm. Define

$$T = W + B = W + D D^T. \quad (12)$$

By convexity, the solution to (9) satisfies the KKT conditions ([Boyd & Vandenberghe, 2004](#), Chapter 5.5). Differentiating (9) with respect to the $(G-1) \times 1$ vector v_j formed by the j th row of V leads to

$$V^t T_j - d_j + \lambda u_j = 0, \quad (13)$$

where T_j is the j th column of matrix T in (12), d_j is a $(G-1) \times 1$ vector formed by the j th row of matrix D in (4) and u_j is the subgradient of $\|v_j\|_2$:

$$u_j = \begin{cases} \frac{v_j}{\|v_j\|_2}, & \text{if } v_j \neq 0; \\ \in \{u : \|u\|_2 \leq 1\}, & \text{if } v_j = 0. \end{cases}$$

Algorithm 1 Block-coordinate descent algorithm.

Given: $k = 1$

$V^{(0)} \leftarrow 0$

repeat

$\bar{V} \leftarrow V^{(k-1)}$

for $j = 1$ **to** p **do**

$v_j^{(k)} \leftarrow \left(1 - \frac{\lambda}{\|d_j - \sum_{i \neq j} t_{ij} \bar{v}_i\|_2}\right)_+ \left(d_j - \sum_{i \neq j} t_{ij} \bar{v}_i\right) / t_{jj}$

end for

$k \leftarrow k + 1$

until $k = k_{\max}$ or $V^{(k)}$ satisfies stopping criterion.

Solving (13) further with respect to v_j leads to $v_j = \left(d_j - \sum_{i \neq j} t_{ij} v_i - \lambda u_j\right) / t_{jj}$, where t_{ij} are the elements of matrix T . This leads to the block-coordinate descent algorithm.

Note that if $\lambda \geq \max_{1 \leq i \leq p} \|d_i\|_2$, then $\hat{V}(\lambda) = 0$. Moreover, if T is non-singular, by applying the vectorization operator (9) can be rewritten as

$$\hat{V}(\lambda) = \arg \min_{V \in \mathbb{R}^{p \times G-1}} \frac{1}{2} \left\| \text{vec}(D^t T^{-1/2}) - (T^{1/2} \otimes I_{G-1}) \text{vec}(V^t) \right\|_2^2 + \lambda \sum_{i=1}^p \|v_i\|_2.$$

This formulation corresponds to a group lasso optimization problem (Yuan & Lin, 2006).

3 Theory

In this section we analyze the variable selection performance of the estimator $\hat{V}(\lambda)$ defined in (9). Denote the support of Ψ by $A = \{j : \|\psi_j\|_2 \neq 0\}$, where ψ_j is the j th row of Ψ , and assume that the support is sparse: $s \ll p$ where $s = \text{card}(A)$. Denote the support of $\hat{V}(\lambda)$ by $\hat{A} = \{j : \|\hat{v}_j(\lambda)\|_2 \neq 0\}$. In Theorem 1 we establish lower bounds on $P(A = \hat{A})$, which we use in Corollary 1 to derive asymptotic conditions under which $P(A = \hat{A}) \rightarrow 1$. In other words, we prove the variable selection consistency of $\hat{V}(\lambda)$.

In Section 2.4 we established an equivalence between our proposal and the proposal of Mai et al. (2012) for the two-group case. This connection allows us to apply theoretical results from Mai et al. (2012) to our method and to extend them to the multi-group case.

Recall $\tilde{\Psi} = \Sigma_W^{-1} \Delta$ from (5) and note that $\{j : \|\tilde{\psi}_j\|_2 \neq 0\} = A$, i.e. Ψ and $\tilde{\Psi}$ have the same support. Let $\Sigma = \Sigma_W + \Sigma_B$ and let $\kappa = \|\Sigma_{A^c A} \Sigma_{AA}^{-1}\|_\infty$, $\phi = \|\Sigma_{AA}^{-1}\|_\infty$, $\tilde{\Psi}_{\min} = \min_{i \in A} \|\tilde{\psi}_i\|_2$ and $\delta = \|\Delta\|_{\infty, 2}$, where Σ_{AA} is the sub-matrix of Σ formed by the intersection of the rows and columns in A .

Theorem 1. Assume $\kappa < 1$ and $(X_i | Y_i = g) \sim N(\mu_g, \Sigma_W)$, $i = 1, \dots, N$. Then

1. For any $\lambda > 0$ and positive $\epsilon \leq \frac{\lambda(1-\kappa)}{(\kappa+1)(\phi\delta+1)+2\phi\lambda}$, $\hat{V}(\lambda)_{A^c} = 0$ with probability at least $1 - t_1$, where

$$t_1 = c_1 p s \exp(-c_2 N s^{-2} \epsilon^2) + 2(G-1)p \exp(-c_3 N \epsilon^2).$$

2. For any $\lambda < \frac{\tilde{\Psi}_{\min}}{\phi}$ and $\epsilon < \frac{\tilde{\Psi}_{\min} - \lambda\phi}{\phi(1+\phi\delta+\tilde{\Psi}_{\min})}$ none of the elements of $\hat{V}(\lambda)_A$ are zero with probability at least $1 - t_2$, where

$$t_2 = c_1 s^2 \exp(-c_2 N s^{-2} \epsilon^2) + 2(G-1)s \exp(-c_3 N \epsilon^2).$$

3. For any positive $\epsilon < \frac{\lambda}{1+\phi\delta+2\phi\lambda}$

$$P\left(\|\hat{V}(\lambda)_A - \tilde{\Psi}\|_{\infty,2} \leq 2\phi\lambda\right) \geq 1 - c_1 s^2 \exp(-c_2 N s^{-2} \epsilon^2) + 2(G-1)s \exp(-c_3 N \epsilon^2).$$

We provide the full proof of Theorem 1 in the Appendix A in order to make the paper self-contained. While the motivation for the proposed optimization problem doesn't rely on the normality assumption, the normality assumption simplifies the proof. We discuss possible extensions to the non-normal case in the Appendix B.

Corollary 1. *From Theorem 1 it follows that if*

1. $N \rightarrow \infty$, $p \rightarrow \infty$, $G = O(1)$ and $\frac{\log(ps)s^2}{N} \rightarrow 0$;
2. $\sqrt{\frac{\log(ps)s^2}{N}} \ll \lambda_N \ll \tilde{\Psi}_{\min}$,

then $P(A = \hat{A}) \rightarrow 1$ and $P(\|\hat{V}(\lambda)_A - \tilde{\Psi}\|_{\infty,2} \leq 2\phi\lambda) \rightarrow 1$.

While preparing this manuscript, we became aware of the improved results on the variable selection consistency of sparse two-group discriminant analysis proposal of Mai et al. (2012). Specifically, Kolar & Liu (2013) show variable selection consistency of $\hat{V}_{DSDA}(\lambda)$ under the condition that $N \geq Cs \log((p-s) \log(N))$ for some constant $C > 0$. Given Proposition 5, these improved rates directly apply to our proposal in the case $G = 2$. We think it may be possible to extend these results to the case $G > 2$, but the details are not straightforward.

4 Simulation Results

In this section we evaluate the performance of the estimator $\hat{V}(\lambda)$ defined in (9) against the alternative methods proposed in the literature. We refer to our proposal as MGSDA for Multi-Group Sparse Discriminant Analysis. We fix sample size n for each group equal to 100 and the number of features p equal to 100 and 800. The test datasets are the same size as the training datasets and are generated independently.

Conditional on the group g , the samples are drawn independently from the multivariate normal distribution $N(\mu_g, \Sigma_W)$. The following structures for Σ_W are considered in all the simulations:

1. **Identity:** $\Sigma_W = I$.
2. **Equicorrelation:** $\Sigma_W = (\sigma_{ij})_{p \times p}$ with $\sigma_{ii} = 1$ and $\sigma_{ij} = 0.5$ for $i \neq j$.
3. **Autoregressive:** $\Sigma_W = (\sigma_{ij})_{p \times p}$ with $\sigma_{ij} = 0.8^{|i-j|}$ for $1 \leq i, j \leq p$.

4. **Bernoulli:** $\Sigma_W = \Omega^{-1}$ with $\Omega = (B + \delta I)/(1 + \delta)$. Here $B = (b_{ij})_{p \times p}$ with $b_{ii} = 1$ for $1 \leq i \leq p$, $b_{ij} = b_{ji} = 0.5 \times \text{Ber}(1, 0.2)$ for $1 \leq i \leq s_0$, $i < j \leq p$ and $b_{ij} = b_{ji} = 0.5$ for $s_0 + 1 \leq i \leq p$, $i < j \leq p$. δ is taken as $\delta = \max(-\lambda_{\min}(B), 0) + 0.05$ to ensure that Ω is positive definite.
5. **Data Based:** $\Sigma_W = (1 - \alpha)S + \alpha I$, where $\alpha = 0.01$ and S is a sample correlation matrix estimated from the most variable $p = 800$ features of Ramaswamy dataset (Ramaswamy et al., 2001). The dataset is available from <http://www-stat.stanford.edu/~tibs/ElemStatLearn/>.

Structures 1-3 have been used in simulation studies in the LDA literature (Cai & Liu, 2011; Witten & Tibshirani, 2011; Mai et al., 2012). The Bernoulli structure was considered by Cai & Liu (2011). We view the Data Based structure as an approximation to a covariance structure that is realistic in practical settings.

4.1 The two-group case


This simulation scenario considers the classification between the two groups with $\mu_1 = 0_p$ and $\mu_2 = (1_s, 0_{p-s})$ for covariance structures 1-4. For covariance structure 5 we take $\mu_2 = (d_s, 0_{p-s})$ with d ranging from 0.1 to 0.5 since in this case the Bayes error is almost zero for $\mu_2 = (1_s, 0_{p-s})$. The simulations are performed for the values of s equal to 10 and 30 for structures 1-4 and s equal to 10 for structure 5.

Mai et al. (2012) perform extensive simulations to compare their proposal with methods of Wu et al. (2009), Witten & Tibshirani (2011), Tibshirani et al. (2003) and Fan & Fan (2008). In all the settings, the method of Mai et al. (2012) performs the best in terms of misclassification error. Given Proposition 5, we do not compare MGSDA with any of these methods. On the other hand, Cai & Liu (2011) also show that their proposal performs the best when compared to Shao et al. (2011), Fan & Fan (2008) and Tibshirani et al. (2003). To our knowledge, no comparison was performed between the methods of Mai et al. (2012) and Cai & Liu (2011), therefore in this section we compare our results to the results of Cai & Liu (2011). We follow the terminology of Cai & Liu (2011) and refer to their method as Linear Programming Discriminant (LPD). We also evaluate the performance of $\tilde{\Psi} = \Sigma_W^{-1} \Delta$. We refer to $\tilde{\Psi}$ as the Oracle.

We note that the LPD requires additional regularization of the within-group sample covariance matrix: $\tilde{W} = W + \rho I$. This regularization is needed to generate a feasible starting point for the optimization algorithm. Cai & Liu (2011) suggest taking $\rho \leq \sqrt{\log p / N}$. In our simulations $N = 200$ and therefore $\rho = 0.15$ satisfies this requirement for both $p = 100$ and $p = 800$. We also try $\rho = 2$ to examine how the choice of ρ affects the misclassification rate.

The misclassification error rates as percentages over 100 replications for covariance structures 1-4 are reported in Table 1. The corresponding number of selected features and the number of false positive features is reported in Table 2. We define the feature j as a false positive if the corresponding component of estimated canonical vector \hat{V} is non-zero, $\hat{v}_j \neq 0$, but $\mu_{1j} - \mu_{2j} = 0$. Note that the population canonical vector Ψ is truly sparse only in the Identity case, it is only approximately sparse in other scenarios. Comparing MGSDA with the best results of LPD show that the methods have similar error rates when the covariance

Figure 1: Mean misclassification error rate in percentage over 25 replications for Data Based covariance structure as a function of difference in means d , $G = 2$.



matrix is Identity or Equicorrelation. MGSDA outperforms LPD for the Autoregressive covariance structure, however LPD performs significantly better for the Bernoulli covariance structure when $p = 800$. The methods select comparable numbers of features in all scenarios.

The mean misclassification rates for the Data Based covariance structure are reported in Figure 1. In this case MGSDA performs significantly better than LPD regardless of the choice of ρ . The difference in misclassification rates is especially noticeable when the difference in means d is small.

The error rates of LPD with $\rho = 0.15$ and $\rho = 2$ are similar for most of the covariance structures, however they are significantly different for Bernoulli structure when $p = 800$ and for the Data Based structure. Table 2 reveals that ρ can also have a significant effect on the number of selected features (the difference is especially noticeable when $s = 30$ and $p = 800$). This suggests that the choice of ρ can significantly affect the performance of the LPD, with smaller values of ρ likely to result in smaller misclassification error. Unfortunately it remains unclear how to choose the optimal ρ in practical settings.

4.2 The multi-group case

This simulation scenario considers the classification between the three groups with $\mu_1 = 0_p$, $\mu_2 = (1_{s/2}, -1_{s/2}, 0_{p-s})$ and $\mu_3 = (-1_{s/2}, 1_{s/2}, 0_{p-s})$. As in the two-group case, we vary the value of μ_2 and μ_3 for the covariance structure 5: $\mu_2 = (d_{s/2}, -d_{s/2}, 0_{p-s})$ and $\mu_3 = (-d_{s/2}, d_{s/2}, 0_{p-s})$ with d ranging from 0.1 to 0.5. The simulations are performed for the values of s equal to 10 and 30 for structures 1-4 and s equal to 10 for structure 5.

The LPD method of Cai & Liu (2011) is developed for the two-group setting. Though it can be generalized to the multi-group case, this generalization is not unique. Among the popular methods are “one versus one” and “one versus all” approaches (Hastie et al., 2009,

Table 1: Mean misclassification error rates as percentages over 100 replications, $G = 2$, standard deviation is given in brackets.

| Covariance | s | p | MGSDA | LPD, $\rho = 0.15$ | LPD, $\rho = 2$ | Oracle |
|-----------------|----|-----|-------------|--------------------|-----------------|-------------|
| Identity | 10 | 100 | 6.65(2.07) | 6.75(2.04) | 6.17(1.94) | 5.58(1.89) |
| | 10 | 800 | 7.32(2.09) | 6.84(1.97) | 6.44(1.73) | 5.75(1.56) |
| | 30 | 100 | 0.9(0.77) | 0.67(0.7) | 0.49(0.53) | 0.4(0.5) |
| | 30 | 800 | 0.83(0.69) | 1.09(0.86) | 0.46(0.5) | 0.32(0.39) |
| Equicorrelation | 10 | 100 | 3.32(1.25) | 3.38(1.7) | 3.02(1.58) | 1.51(0.89) |
| | 10 | 800 | 3.11(1.25) | 2.98(1.39) | 2.79(1.17) | 1.45(0.81) |
| | 30 | 100 | 0.55(0.53) | 0.55(0.67) | 0.52(0.73) | 0.06(0.2) |
| | 30 | 800 | 0.27(0.38) | 0.5(0.61) | 0.56(0.77) | 0(0) |
| Autoregressive | 10 | 100 | 19.02(2.91) | 20.83(3.17) | 23.88(2.99) | 16.65(2.48) |
| | 10 | 800 | 22.29(3.26) | 23.59(3.35) | 24.5(3.06) | 16.05(2.59) |
| | 30 | 100 | 13.72(2.68) | 15.26(2.91) | 15.85(2.84) | 10.97(2.14) |
| | 30 | 800 | 16.57(2.71) | 17.23(3.26) | 16.81(2.64) | 11.13(1.95) |
| Bernoulli | 10 | 100 | 6.12(1.69) | 5.88(1.48) | 5.75(1.62) | 4.37(1.35) |
| | 10 | 800 | 37.14(6.04) | 17.03(3.4) | 28.62(3.59) | 4.6(1.49) |
| | 30 | 100 | 0.35(0.42) | 0.3(0.38) | 0.22(0.34) | 0.05(0.15) |
| | 30 | 800 | 8.27(2.81) | 3.29(1.4) | 7.17(2.74) | 0.04(0.14) |


Table 2: Mean number of selected features and false positive features over 100 replications, $G = 2$, standard deviation is given in brackets.

| Covariance | s | p | All Features | | | False Positives | | |
|-----------------|----|-----|--------------|--------------------|-----------------|-----------------|--------------------|-----------------|
| | | | MGSDA | LPD, $\rho = 0.15$ | LPD, $\rho = 2$ | MGSDA | LPD, $\rho = 0.15$ | LPD, $\rho = 2$ |
| Identity | 10 | 100 | 20(7) | 19(14) | 19(15) | 10(7) | 9(14) | 9(15) |
| | 10 | 800 | 29(16) | 25(26) | 24(29) | 19(16) | 15(26) | 14(29) |
| | 30 | 100 | 40(7) | 38(11) | 34(13) | 11(7) | 9(11) | 4(13) |
| | 30 | 800 | 51(15) | 125(93) | 56(83) | 22(15) | 95(93) | 26(83) |
| Equicorrelation | 10 | 100 | 51(5) | 55(12) | 73(10) | 41(5) | 45(12) | 63(10) |
| | 10 | 800 | 84(13) | 90(54) | 128(48) | 74(13) | 80(54) | 118(48) |
| | 30 | 100 | 77(4) | 78(9) | 93(6) | 49(4) | 49(8) | 63(6) |
| | 30 | 800 | 147(13) | 112(54) | 177(40) | 119(13) | 84(54) | 147(40) |
| Autoregressive | 10 | 100 | 19(6) | 20(14) | 30(21) | 14(6) | 13(13) | 20(21) |
| | 10 | 800 | 32(15) | 23(21) | 27(31) | 26(15) | 16(20) | 17(31) |
| | 30 | 100 | 26(7) | 31(15) | 46(17) | 12(6) | 14(13) | 17(16) |
| | 30 | 800 | 41(20) | 41(51) | 76(92) | 28(19) | 25(49) | 48(91) |
| Bernoulli | 10 | 100 | 24(9) | 18(11) | 20(18) | 14(9) | 8(11) | 10(18) |
| | 10 | 800 | 43(33) | 70(83) | 19(14) | 38(31) | 60(83) | 9(14) |
| | 30 | 100 | 43(8) | 39(14) | 33(11) | 14(8) | 9(14) | 3(11) |
| | 30 | 800 | 116(32) | 216(117) | 43(18) | 90(31) | 187(116) | 13(18) |

p. 658). In addition to requiring the computation of a larger number of discriminant vectors ($G(G - 1)/2$ and G correspondingly), these approaches can disagree in their classification rules as well as in selected features. Given this ambiguity, we do not compare our method to the LPD in the multi-group case.

We were able to find only two methods in the literature that specifically consider sparse discriminant analysis in the multi-group case: penalizedLDA by [Witten & Tibshirani \(2011\)](#) and sparseLDA by [Clemmensen et al. \(2011\)](#). Both methods find canonical vectors sequentially and are nonconvex. We compare their performance with MGSDA and also evaluate the performance of $\tilde{\Psi} = \Sigma_W^{-1} \Delta$. Again, we refer to $\tilde{\Psi}$ as the Oracle.

Figure 2: Mean misclassification error rate in percentage over 25 replications for Data Based covariance structure as a function of difference in means d , $G = 3$.



The mean misclassification error rates as percentages over 100 replications for each combination of parameters are reported in Table 3. The number of selected features and the number of false positive features is reported in Table 4. Similar to the two-group case, we define the feature j as a false positive if the corresponding row of estimated canonical vector matrix \hat{V} is non-zero, $\|\hat{v}_j\|_2 \neq 0$, but $\mu_{1j} = \mu_{2j} = \mu_{3j}$. Also as before, the population canonical vectors matrix Ψ is truly row sparse only in the Identity case, it is only approximately row sparse in other scenarios. The results suggest that all three methods are comparable in terms of misclassification rate except for the Autoregressive covariance structure. In this scenario, both MGSDA and sparseLDA outperform the penalizedLDA. In terms of the number of features, MGSDA tends to select fewer than the competitors. Hence, MGSDA achieves the best tradeoff between the misclassification error and sparsity of the solution.

The mean misclassification rates for the Data Based covariance structure are reported in Figure 2. It can be seen that the penalizedLDA performs significantly worse than both MGSDA and sparseLDA. This is not a surprising result since the Data Based covariance structure is far from diagonal, which is an underlying assumption of penalizedLDA.

4.3 Implementation Details

The method of Cai & Liu (2011) is implemented using `linprogPD` function from the package CLIME from CRAN. Note that `linprogPD` almost never returns a sparse solution. However, all the values below the precision level should be treated as zeroes (Cai et al., 2011). We used the default value of 10^{-3} for precision. The grid for the tuning parameter is chosen from 0.01 to 0.5 by 0.01. The method of Witten & Tibshirani (2011) is implemented using the package `penalizedLDA` from CRAN. The grid for tuning parameter is chosen from 0 to 1 by 0.01. The method of Clemmensen et al. (2011) is implemented using the package `sparseLDA`

Table 3: Mean misclassification error rates as percentages over 100 replications, $G = 3$, standard deviation is given in brackets.

| Covariance | s | p | MGSDA | penalizedLDA | sparseLDA | Oracle |
|-----------------|----|-----|-------------|--------------|-------------|------------|
| Identity | 10 | 100 | 9.11(1.52) | 8.4(1.6) | 9.39(1.68) | 7.83(1.41) |
| | 10 | 800 | 9.22(1.73) | 7.92(1.5) | 9.58(1.87) | 7.67(1.43) |
| | 30 | 100 | 1.06(0.67) | 0.5(0.42) | 0.93(0.63) | 0.42(0.37) |
| | 30 | 800 | 1.43(0.81) | 0.52(0.39) | 1.14(0.72) | 0.43(0.35) |
| Equicorrelation | 10 | 100 | 2.15(0.97) | 2.34(1.18) | 2.03(0.94) | 1.65(0.86) |
| | 10 | 800 | 2.19(0.89) | 2.12(1.1) | 2.15(0.85) | 1.68(0.84) |
| | 30 | 100 | 0.23(0.38) | 0.3(1.11) | 0.26(0.44) | 0.01(0.05) |
| | 30 | 800 | 0.31(0.43) | 0.04(0.11) | 0.39(0.52) | 0.01(0.06) |
| Autoregressive | 10 | 100 | 6.83(1.4) | 16.87(2.16) | 6.34(1.32) | 4.87(1) |
| | 10 | 800 | 7.29(1.77) | 16.53(2.44) | 7.47(2.54) | 4.9(1.17) |
| | 30 | 100 | 5.45(1.48) | 16(1.95) | 4.86(1.44) | 3.57(1.05) |
| | 30 | 800 | 5.89(1.53) | 15.46(2.33) | 5.98(2.03) | 3.65(1.05) |
| Bernoulli | 10 | 100 | 11.15(1.91) | 10.56(1.73) | 11.35(1.88) | 8.51(1.62) |
| | 10 | 800 | 43.13(3.06) | 44.84(4) | 41.72(3.19) | 8.56(1.67) |
| | 30 | 100 | 1.54(0.72) | 1.05(0.56) | 1.37(0.74) | 0.59(0.47) |
| | 30 | 800 | 20.59(3.08) | 22.66(4.11) | 16.42(2.57) | 0.63(0.46) |

Table 4: Mean number of selected features and false positive features over 100 replications, $G = 3$, standard deviation is given in brackets.

| Covariance | s | p | All Features | | | False Positives | | |
|-----------------|----|-----|--------------|--------------|-----------|-----------------|--------------|-----------|
| | | | MGSDA | penalizedLDA | sparseLDA | MGSDA | penalizedLDA | sparseLDA |
| Identity | 10 | 100 | 13(7) | 15(15) | 26(11) | 3(7) | 5(15) | 16(11) |
| | 10 | 800 | 11(2) | 15(8) | 24(9) | 1(2) | 5(8) | 14(9) |
| | 30 | 100 | 46(18) | 61(19) | 58(11) | 16(18) | 31(19) | 29(11) |
| | 30 | 800 | 37(11) | 51(93) | 67(18) | 8(10) | 21(93) | 37(18) |
| Equicorrelation | 10 | 100 | 14(8) | 10(1) | 27(12) | 4(8) | 0(1) | 17(12) |
| | 10 | 800 | 12(6) | 12(3) | 28(19) | 2(6) | 2(3) | 18(19) |
| | 30 | 100 | 29(11) | 38(12) | 47(8) | 2(10) | 8(12) | 21(6) |
| | 30 | 800 | 29(14) | 30(0) | 50(7) | 4(13) | 0(0) | 25(3) |
| Autoregressive | 10 | 100 | 21(16) | 12(6) | 27(12) | 13(16) | 2(6) | 20(12) |
| | 10 | 800 | 7(3) | 15(11) | 29(19) | 1(2) | 5(11) | 22(19) |
| | 30 | 100 | 28(16) | 54(25) | 39(11) | 13(14) | 24(25) | 22(10) |
| | 30 | 800 | 14(5) | 36(40) | 49(24) | 2(4) | 6(40) | 34(23) |
| Bernoulli | 10 | 100 | 14(10) | 16(15) | 30(13) | 4(10) | 6(15) | 20(13) |
| | 10 | 800 | 118(115) | 33(61) | 100(46) | 110(114) | 24(61) | 92(45) |
| | 30 | 100 | 51(21) | 59(22) | 61(11) | 22(21) | 29(22) | 31(11) |
| | 30 | 800 | 42(33) | 48(14) | 108(33) | 21(31) | 18(14) | 83(32) |

from CRAN. Each canonical vector is constrained to have between 3 and 80 features. This is quite a restrictive range for tuning, however the **sparseLDA** package produced errors when we used a wider range of features. MGSDA is implemented using the R package **MGSDA**, which is currently under development. The grid for the tuning parameter $\lambda_1 \leq \dots \leq \lambda_{max}$ is chosen adaptively for each dataset with $\lambda_{max} = \max_j \|d_j\|_2$, which according to Section 2.5 corresponds to zero selected features. For each $\lambda_l < \lambda_{max}$ we set $V^{(0)} = \hat{V}(\lambda_{l+1})$. For all the methods, the final tuning parameter is chosen from the respective grid through 5-fold cross-validation to minimize the error rate.

Witten&Tibshirani’s penalizedLDA has significantly faster running time than all other

methods since penalizedLDA assumes that the covariance matrix has diagonal structure. This assumption results in a simplified optimization algorithm, for details we refer to [Witten & Tibshirani \(2011\)](#). The running time of penalizedLDA is followed by MGSDA and sparseLDA. Surprisingly, LPD has the slowest performance. We suspect that this is not due to the method itself, but due to the use of `linprogPD` function in its implementation. A different linear program solver is likely to result in much faster running time, however the use of a general solver makes the method implementation less straightforward.

5 Real Data

5.1 Metabolomics Dataset

Metabolomics is the global study of all metabolites in a biological system under a given set of conditions. Metabolites are the final products of enzymes and enzyme networks whose substrates and products often cannot be deduced from genetic information and whose levels reflect the integrated product of the genome, proteome and environment. Metabolomic readouts thus represent the most direct (or phenotypic) readout of a cells physiologic state. From a technical standpoint, analytical studies of metabolism have been historically limited to one or a limited set of metabolites. However, advances in liquid chromatography and mass spectrometry have recently made it possible to measure hundreds of metabolites and with enough biomass well over 1000, in parallel. Such technologies have thus opened the door to obtaining global biochemical readouts of a cells physiologic state and response to perturbation. Cornell researchers have developed and applied a state-of-the-art metabolomic platform to track the intrabacterial pharmacokinetic fates and pharmacodynamic actions of a given compound within *Mycobacterium tuberculosis* ([Pethe et al., 2010](#); [de Carvalho et al., 2010, 2011](#); [Chakraborty et al., 2013](#)). These studies demonstrate the highly unpredictable nature and identities of these properties even for well-studied antibiotics.

We investigate a (currently unpublished) metabolomics data obtained from Dr. Kyu Rhee, Cornell University, which seeks to systematically elucidate the intrabacterial pharmacokinetic and pharmacodynamic fates and actions of antimycobacterial hit or lead compound series identified in high throughput screens against replicating and non- or slowly replicating forms of *Mycobacterium tuberculosis*.

The data contains measurements of 171 metabolic responses of 68 patients to 25 antibiotics that are administered at different dosage levels. Each measurement is an average of three replicates, normalized to the vehicle control and log2 transformed. 14 out of 25 antibiotics can be divided into the following 5 groups: STREP_AMI(strep, ami), FLQ(lev, moxi), DHFR(nitd2, sri8210, sri 8710, sri 8857), DHPS(smx, snl, aps) and InhA(eta, isoxyl, gsk93). These antibiotics are administered to 35 patients out of 68. In the subsequent analysis we only focus on these 5 groups of antibiotics and do not consider the dosage levels.

We compare the performance of our method, penalizedLDA ([Witten & Tibshirani, 2011](#)) and sparseLDA ([Clemmensen et al., 2011](#)) on this data set using the following measures: mean 5-fold cross-validation error and the number of selected features over 20 replications. We do not perform random splits into the training and test set due to the small sample size. The results are reported in Table 5.

Table 5: Mean cross-validation error and mean number of selected features over 20 replications on metabolomics dataset, standard deviation is given in brackets.

| | MGSDA | penalizedLDA | sparseLDA |
|----------|------------|--------------|------------|
| CV error | 0.15(0.08) | 0.09(0.15) | 0.02(0.06) |
| Features | 24.75(15) | 158.75(18.4) | 30.2(8.6) |

The results show that all three methods perform very well in terms of cross-validation error making on average less than one mistake. Such a good performance suggests that there is a significant difference in the metabolic responses between the 5 groups of antibiotics. However, penalizedLDA achieves this performance by selecting almost all of the metabolites, whereas our method and sparseLDA use less than 20% of the original features. Note that there is a substantial variation between the replications due to the small sample size of the data.

We further estimate four canonical vectors using MGSDA with $\lambda = 0.57$ and illustrate the projected data in Figure 3. Note that 8 selected metabolites provide perfect linear separation between the groups. $\lambda = 0.57$ is chosen as one of the twenty tuning parameters from above replications of cross-validation splits. We have tried the other values of λ as well, however they all provided perfect linear separation between the groups with projected data being very similar to Figure 3. Though there is a variation between the cross-validation replications due to the small sample size of the data, this variation has negligible effect on the final projection.

5.2 14 Cancer Dataset

In this section we compare the performance of MGSDA, penalizedLDA (Witten & Tibshirani, 2011) and sparseLDA (Clemmensen et al., 2011) on 14 cancer dataset by Ramaswamy et al. (2001). This dataset contains 16063 gene expression measurements collected on 198 samples. Each sample belongs to one of the 14 cancer classes. The data set can be obtained from <http://statweb.stanford.edu/~tibs/ElemStatLearn/>. We selected this data set as it is publicly available and has been previously analyzed by a number of authors including Witten & Tibshirani (2011).

Following the recommendation of Hastie et al. (2009, p. 654), we first standardize the data to have mean zero and standard deviation one for each patient. Next we restrict the analysis to 3000 genes with largest standard deviation among the samples to reduce the overall computational cost. Following the approach taken by Witten & Tibshirani (2011), we perform 10 independent splits of the data set into the training set containing 75% of the samples and the test set containing 25% of the samples. The tuning parameter for all methods is selected using 5-fold cross-validation on the training set. The mean misclassification error on the test set and the mean number of selected features over 10 splits are reported in Table 6. MGSDA and sparseLDA perform better than penalizedLDA in terms of the misclassification error and select much smaller number of features. The performance of our method and sparseLDA is comparable on this dataset.

Figure 3: Metabolomics dataset projected onto 4 column vectors of V , $k = 8$ metabolite features are used.




Table 6: Mean number of misclassified samples and mean number of selected features over 10 splits on 14 cancer dataset, standard deviation is given in brackets.

| | MGSDA | penalizedLDA | sparseLDA |
|----------|---------------|--------------|--------------|
| Error | 9.9(2.13) | 15.4(2.95) | 10.7(2.06) |
| Features | 238.3(135.16) | 2845(109.46) | 268.3(84.63) |

6 Discussion

In this paper we introduce a novel algorithm that estimates population canonical vectors in the multi-group setting. The proposed method is a natural generalization of the two-group methods that were previously studied in the literature. In addition to being computationally tractable, the method performs feature selection which results in sparse canonical vectors. The group penalty eliminates the features from all canonical vectors at once with the remaining non-zero features being the same for all the vectors.

One possible extension of the proposed method is to allow canonical vectors to have different sparsity patterns. This goal can be achieved through the addition of the within-row penalty term to the objective function (9). Such an estimation procedure has been already considered in the regression context; for example, [Simon et al. \(2013\)](#) propose the following

optimization problem:

$$\hat{\beta}(\lambda, \alpha) = \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|Y - X\beta\|_2^2 + (1 - \alpha)\lambda \sum_{g=1}^G \sqrt{p_g} \|\beta^{(g)}\|_2 + \alpha\lambda \|\beta\|_1,$$

where p_g is the size of group g . In our case this approach results in

$$\hat{V}(\lambda, \alpha) = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \text{Tr}(V^t W V) + \frac{1}{2} \|D^t V - I\|_F^2 + (1 - \alpha)\lambda \sum_{i=1}^p \|v_i\|_2 + \alpha\lambda \|V\|_1.$$

Another possible extension is to perform canonical vectors selection in addition to feature selection, which will enhance the interpretability, especially when the number of groups G is large. This goal can be achieved through the addition of the nuclear norm penalty term to the objective function (9):

$$\hat{V}(\lambda, \alpha) = \arg \min_{V \in \mathbb{R}^{p \times (G-1)}} \frac{1}{2} \text{Tr}(V^t W V) + \frac{1}{2} \|D^t V - I\|_F^2 + \lambda \sum_{i=1}^p \|v_i\|_2 + \alpha \|V\|_*.$$

Depending on the value of $\alpha > 0$, the resulting matrix \hat{V} has rank that is less than $G - 1$, effectively resulting in a lower-dimensional eigenspace.

Both extensions result in convex optimization problems, but require additional modifications to optimization algorithm 1. An interesting direction for future research is to examine how these extensions compare to the original method in different scenarios.

Appendix A

Proof of Proposition 1

Proof. The proof is only given for matrix B , the proof for matrix Σ_B is similar.

1. Consider the equal group case: $n_1 = \dots = n_G = n$ and $N = Gn$. It follows that $\bar{X} = \sum_{i=1}^G \bar{X}_G / G$ and therefore $B = \frac{1}{N} \sum_{g=1}^G n(\bar{X}_g - \bar{X})(\bar{X}_g - \bar{X})^t = \frac{1}{N} X^t \left\{ \frac{1}{\sqrt{n}} \mathbf{1}_g \right\} C \left\{ \frac{1}{\sqrt{n}} \mathbf{1}_g \right\}^t X$, where C is the centering matrix and $\left\{ \frac{1}{\sqrt{n}} \mathbf{1}_g \right\}$ is a $N \times G$ matrix formed by G columns $\frac{1}{\sqrt{n}} \mathbf{1}_g$ such that $(\mathbf{1}_g)_j = 1$ if j th observation belongs to the g th group and $(\mathbf{1}_g)_j = 0$ otherwise. Note that $C = H^t H$ where H is the Helmert matrix of size G with its first row removed (Searle, 2006). Therefore, B can be rewritten as $B = \frac{1}{N} X^t \left\{ \frac{1}{\sqrt{n}} \mathbf{1}_g \right\} H^t H \left\{ \frac{1}{\sqrt{n}} \mathbf{1}_g \right\}^t X = D D^t$, where $D = \frac{1}{\sqrt{N}} X^t \left\{ \frac{1}{\sqrt{n}} \mathbf{1}_g \right\} H^t$.

2. Consider the general case where each group has size n_i . Similar to the equal group case, B can be rewritten as $B = \frac{1}{N} X^t \left\{ \frac{1}{\sqrt{n_i}} \mathbf{1}_i \right\} \tilde{C} \left\{ \frac{1}{\sqrt{n_i}} \mathbf{1}_i \right\}^t X$, where $\tilde{C} = I_G - \frac{1}{\sqrt{N}} K K^t$ and $K = (\sqrt{n_1} \dots \sqrt{n_G})^t$. Next we show that similar to C , \tilde{C} can be decomposed as $\tilde{C} = \tilde{H}^t \tilde{H}$ and \tilde{H} is a $G - 1 \times G$ adjusted Helmert matrix. Recall that \tilde{H} satisfies

$$I_G - \frac{1}{\sqrt{N}} K K^t = \tilde{H}^t \tilde{H}.$$

In other words, (K, \tilde{H}^t) is an orthogonal matrix. The $G - 1$ orthogonal contrasts for unbalanced data have the following form (Searle, 2006, p 51):

$$\delta_r = \sqrt{n_{r+1}} \left(\sum_{h=1}^r n_h (\bar{X}_h - \bar{X}_{r+1}) \right).$$

Denote by h_r the rows of \tilde{H} . Then it follows that

$$h_r \left\{ \frac{1}{\sqrt{n_i}} \mathbf{1}_i \right\}^t X = C_r \delta_r,$$

where C_r is some constant. This means that $h_{rj} = C_r \sqrt{n_{r+1} n_j}$ for $j = 1, \dots, r$; $h_{r(r+1)} = -C_r \sum_{i=1}^r n_i$ and $h_{rj} = 0$ for $j > (r+1)$. To find C_r , we use the fact that $h_r h_r^t = 1$. Let $s_r = \sum_{i=1}^r n_i$. Then C_r satisfies:

$$\begin{aligned} C_r^2 \left(\sum_{j=1}^r n_{r+1} n_j + s_r^2 \right) &= 1; \\ C_r^2 (n_{r+1} s_r + s_r^2) &= 1; \\ C_r^2 s_{r+1} s_r &= 1. \end{aligned}$$

From the last equation $C_r = \frac{1}{\sqrt{s_{r+1} s_r}}$. Combining the results it follows that

$$B = \frac{1}{N} X^t \left\{ \frac{1}{\sqrt{n_g}} \mathbf{1}_i \right\} \tilde{H} \tilde{H}^t \left\{ \frac{1}{\sqrt{n_g}} \mathbf{1}_i \right\}^t X = D D^t,$$

where $D = \frac{1}{\sqrt{N}} X^t \left\{ \frac{1}{\sqrt{n_g}} \mathbf{1}_i \right\} \tilde{H}^t$ and $D_r = \frac{1}{\sqrt{N}} C_r \delta_r = \frac{\sqrt{n_{r+1}} (\sum_{h=1}^r n_h (\bar{X}_h - \bar{X}_{r+1}))}{\sqrt{N s_{r+1} s_r}}$. □

Proof of Proposition 2

Proof. Denote $\Psi = \Sigma_W^{-1} \Delta P$, where P is an orthogonal matrix such that $\Delta^t \Sigma_W^{-1} \Delta = P \Lambda P^t$. It follows that $\Sigma_W^{-1} \Sigma_B \Psi = \Sigma_W^{-1} \Delta \Delta^t \Sigma_W^{-1} \Delta P = \Sigma_W^{-1} \Delta P \Lambda = \Psi \Lambda$. Hence, Ψ is the matrix of eigenvectors of $\Sigma_W^{-1} \Sigma_B$. The proof for V is analogous. □

Proof of Proposition 3

Proof. The proof is only given for the sample classification rule $\hat{h}_V(x)$, the proof for the population classification rule $h_\Psi(x)$ is analogous. Define $Z = XV$. Using V , a new observation $x \in \mathbb{R}^p$ is classified to group $\hat{h}_V(x)$, where

$$\hat{h}_V(x) = \arg \min_{1 \leq j \leq G} (V^t x - \bar{Z}_j)^t (V^t W V)^{-1} (V^t x - \bar{Z}_j) - 2 \log \frac{n_j}{N}.$$

Consider a new classification rule $\hat{h}_{V'}(x)$ based on $V' = VR$ with $R \in \mathbb{O}^{G-1}$. Then

$$Z' = XV' = XVR = ZR$$

and

$$\begin{aligned}
\hat{h}_{V'}(x) &= \arg \min_{1 \leq j \leq G} (V'^t x - \bar{Z}'_j)^t (V'^t W V')^{-1} (V'^t x - \bar{Z}'_j) - 2 \log \frac{n_j}{N} \\
&= \arg \min_{1 \leq j \leq G} (R^t V^t x - R^t \bar{Z}_j)^t (R^t V^t W V R)^{-1} (R^t V^t x - R^t \bar{Z}_j) - 2 \log \frac{n_j}{N} \\
&= \arg \min_{1 \leq j \leq G} (V^t x - \bar{Z}_j)^t R R^{-1} (V^t W V)^{-1} (R^t)^{-1} R^t (V^t x - \bar{Z}_j) - 2 \log \frac{n_j}{N} \\
&= \hat{h}_V(x).
\end{aligned}$$

□

Proof of Proposition 4

Proof. From the definition of Υ_ρ , $(\Sigma_W + \rho \Sigma_B)^{-1} \rho \Sigma_B \Upsilon_\rho = \Upsilon_\rho \Lambda$. It follows that

$$\begin{aligned}
\rho \Sigma_B \Upsilon_\rho &= \Sigma_W \Upsilon_\rho \Lambda + \rho \Sigma_B \Upsilon_\rho \Lambda; \\
\rho \Sigma_B \Upsilon_\rho (I - \Lambda) &= \Sigma_W \Upsilon_\rho \Lambda; \\
\Sigma_W^{-1} \Sigma_B \Upsilon_\rho &= \Upsilon_\rho \frac{1}{\rho} \Lambda (I - \Lambda)^{-1}.
\end{aligned}$$

From the last equation it follows that Υ_ρ is the matrix of eigenvectors of $\Sigma_W^{-1} \Sigma_B$. Since the eigenvectors are unique up to normalization, the statement of the proposition follows.

□

Proof of Proposition 5

Proof. By definition $D = \frac{\sqrt{n_1 n_2}}{N} (\bar{X}_1 - \bar{X}_2)$. Therefore

$$\begin{aligned}
\hat{V}_{DSDA}(\lambda) &= \arg \min_{V \in \mathbb{R}^p} \frac{1}{2} V^t (W + D D^t) V - \frac{N}{\sqrt{n_1 n_2}} D^t V + \lambda \|V\|_1 \\
&= \arg \min_{V \in \mathbb{R}^p} \frac{1}{2} \frac{\sqrt{n_1 n_2}}{N} V^t (W + D D^t) V - D^t V + \frac{\lambda \sqrt{n_1 n_2}}{N} \|V\|_1.
\end{aligned}$$

Define

$$f_{DSDA}(V, \lambda) = \frac{1}{2} \frac{\sqrt{n_1 n_2}}{N} V^t (W + D D^t) V - D^t V + \frac{\lambda \sqrt{n_1 n_2}}{N} \|V\|_1.$$

Similarly, $\hat{V}(\lambda) = \arg \min_{V \in \mathbb{R}^p} f(V, \lambda)$, where

$$f(V, \lambda) = \frac{1}{2} V^t (W + D D^t) V - D^t V + \lambda \|V\|_1.$$

Note that

$$\begin{aligned}
f\left(\frac{\sqrt{n_1 n_2}}{N} V, \lambda\right) &= \frac{\sqrt{n_1 n_2}}{N} \left(\frac{1}{2} \frac{\sqrt{n_1 n_2}}{N} V^t (W + D D^t) V - D^t V + \lambda \|V\|_1\right) \\
&= \frac{\sqrt{n_1 n_2}}{N} f_{DSDA}\left(V, \frac{N}{\sqrt{n_1 n_2}} \lambda\right).
\end{aligned}$$

It follows that $\hat{V}(\lambda) = \frac{N}{\sqrt{n_1 n_2}} \hat{V}_{DSDA}\left(\frac{N}{\sqrt{n_1 n_2}} \lambda\right)$.

□

Auxillary lemmas for Theorem 1.

Lemma 1. $\|AB\|_{\infty,2} \leq \|A\|_{\infty}\|B\|_{\infty,2}$

Proof. This inequality is a special case of Lemma 8 in [Obozinski et al. \(2011\)](#). Note that

$$\|A\|_{\infty,2} = \max_i \|a_i\|_2 = \max_i \max_{\|y_i\|_2 \leq 1} |y^t a_i| = \max_{\|y\|_2 \leq 1} \max_i |y^t a_i| = \max_{\|y\|_2 \leq 1} \|Ay\|_2.$$

It follows that

$$\|AB\|_{\infty,2} = \max_{\|y\|_2 \leq 1} \|AB y\|_{\infty} \leq \max_{\|y\|_2 \leq 1} \|A\|_{\infty} \|B y\|_{\infty} = \|A\|_{\infty} \max_{\|y\|_2 \leq 1} \|B y\|_{\infty} = \|A\|_{\infty} \|B\|_{\infty,2}.$$

□

Lemma 2. Define $F = D - \Delta$. Then there exists constant $c_3 > 0$ such that

$$P(\|F\|_{\infty,2} \geq \epsilon) \leq 2p(G-1) \exp(-c_3 N \epsilon^2).$$

Proof. This lemma is a multi-group generalization of Lemma A1 in [Mai et al. \(2012\)](#). From definition of Δ , its r th column has the form

$$\Delta_r = \frac{\sqrt{\pi_{r+1}} (\sum_{i=1}^r \pi_i (\mu_i - \mu_{r+1}))}{\sqrt{\sum_{i=1}^r \pi_i} \sqrt{\sum_{i=1}^{r+1} \pi_i}} = \frac{\sum_{i=1}^r (\mu_i - \mu_{r+1})}{\sqrt{Gr(r+1)}}.$$

From definition of D , its r th column has the form

$$D_r = \frac{\sqrt{n_{r+1}} (\sum_{i=1}^r n_i (\bar{X}_i - \bar{X}_{r+1}))}{\sqrt{N} \sqrt{\sum_{i=1}^r n_i} \sqrt{\sum_{i=1}^{r+1} n_i}} = \frac{\sum_{i=1}^r (\bar{X}_i - \bar{X}_{r+1})}{\sqrt{Gr(r+1)}}.$$

Therefore

$$F_r = \frac{1}{\sqrt{Gr(r+1)}} \sum_{i=1}^r ((\bar{X}_i - \bar{X}_{r+1}) - (\mu_i - \mu_{r+1})).$$

Since the groups are independent and $(\bar{X}_g)_j \sim N((\mu_g)_j, \frac{\sigma_j^2}{n})$ for all $g \in \{1, \dots, G\}$ and $j \in \{1, \dots, p\}$, then for all r :

$$\sum_{i=1}^r (\bar{X}_i - \bar{X}_{r+1})_j \sim N\left(\sum_{i=1}^r (\mu_i - \mu_{r+1})_j, \frac{r(r+1)\sigma_j^2}{n}\right),$$

or equivalently

$$d_{jr} \sim N\left(\delta_{jr}, \frac{\sigma_j^2}{Gn}\right),$$

where d_{jr} are the elements of matrix D and δ_{jr} are the elements of matrix Δ . It follows that for all $r \in \{1, \dots, G-1\}$ and for all $j \in \{1, \dots, p\}$

$$P(|f_{jr}| \geq \epsilon) = P(|d_{jr} - \delta_{jr}| \geq \epsilon) \leq 2 \exp\left(-\frac{N \epsilon^2}{2 \sigma_j^2}\right) \leq 2 \exp(-c N \epsilon^2).$$

Therefore

$$\begin{aligned}
P(\|f_j\|_2 \geq \epsilon) &= P\left(\sqrt{f_{j1}^2 + \dots + f_{(G-1)j}^2} \geq \epsilon\right) \leq P\left(\sqrt{G-1} \max_r |f_{jr}| \geq \epsilon\right) \\
&\leq P\left(\cup_r \left\{|f_{jr}| \geq \frac{\epsilon}{\sqrt{G-1}}\right\}\right) \leq (G-1)P\left(|f_{jr}| \geq \frac{\epsilon}{\sqrt{G-1}}\right) \\
&\leq 2(G-1)\exp(-c_3 N \epsilon^2).
\end{aligned}$$

Applying the union bound over $j \in \{1, \dots, p\}$ gives

$$P(\|F\|_{\infty,2} \geq \epsilon) \leq 2p(G-1)\exp(-c_3 N \epsilon^2).$$

□

Lemma 3. *Let $T = W + B$. Then there exist constants $c_1 > 0$ and $c_2 > 0$ such that*

$$\begin{aligned}
P(\|T_{AA} - \Sigma_{TAA}\|_{\infty} \geq \epsilon) &\leq c_1 s^2 \exp(-c_2 N s^{-2} \epsilon^2); \\
P(\|T_{A^c A} - \Sigma_{T A^c A}\|_{\infty} \geq \epsilon) &\leq c_1 s(p-s) \exp(-c_2 N s^{-2} \epsilon^2).
\end{aligned}$$

Proof. This lemma is a multi-group version of Lemma A.1 in [Mai et al. \(2012\)](#). First, we show that

$$P(|\Sigma_{Tij} - T_{ij}| > \epsilon) \leq c_1 \exp(-c_2 N \epsilon^2).$$

By definition of Σ_T and T :

$$\begin{aligned}
\Sigma_{Tij} - T_{ij} &= \Sigma_{Wij} + \Sigma_{Bij} - \frac{1}{N} \sum_{k=1}^N (X_{ki} - \bar{X}_i)(X_{kj} - \bar{X}_j) \\
&= \Sigma_{Wij} + \sum_{g=1}^G \pi_g (\mu_{gi} - \mu_i)(\mu_{gj} - \mu_j) - \frac{1}{N} \sum_{k=1}^N X_{ki} X_{kj} + \bar{X}_i \bar{X}_j \\
&= \Sigma_{Wij} + \sum_{g=1}^G \pi_g \mu_{gi} \mu_{gj} + \bar{X}_i \bar{X}_j - \mu_i \mu_j - \frac{1}{N} \sum_{g=1}^G \sum_{k \in I_g} X_{ki} X_{kj}.
\end{aligned}$$

Furthermore

$$\frac{1}{n_g} \sum_{k \in I_g} X_{ki} X_{kj} = \frac{1}{n_g} \sum_{k \in I_g} (X_{ki} - \mu_{gi})(X_{kj} - \mu_{gj}) + \mu_{gi}(\bar{X}_{gj} - \mu_{gj}) + \mu_{gj}(\bar{X}_{gi} - \mu_{gi}) + \mu_{gj} \mu_{gi}.$$

Therefore

$$\begin{aligned}
\Sigma_{Tij} - T_{ij} &= \Sigma_{Wij} + \sum_{g=1}^G \pi_g \mu_{gi} \mu_{gj} + \bar{X}_i \bar{X}_j - \mu_i \mu_j - \\
&\quad - \frac{1}{N} \sum_{g=1}^G n_g \left(\frac{1}{n_g} \sum_{k \in I_g} (X_{ki} - \mu_{gi})(X_{kj} - \mu_{gj}) + \mu_{gi}(\bar{X}_{gj} - \mu_{gj}) + \mu_{gj}(\bar{X}_{gi} - \mu_{gi}) + \mu_{gj} \mu_{gi} \right) \\
&= \sum_{g=1}^G \frac{n_g}{N} \left(\Sigma_{Wij} - \frac{1}{n_g} \sum_{k \in I_g} (X_{ki} - \mu_{gi})(X_{kj} - \mu_{gj}) \right) \\
&\quad + \sum_{g=1}^G \frac{n_g}{N} (\mu_{gi}(\mu_{gj} - \bar{X}_{gj}) + \mu_{gj}(\mu_{gi} - \bar{X}_{gi})) + \sum_{g=1}^G \left(\pi_g - \frac{n_g}{N} \right) \mu_{gi} \mu_{gj} + (\bar{X}_i \bar{X}_j - \mu_i \mu_j)
\end{aligned}$$

Under the assumption $\pi_g = \frac{1}{G}$ and $n_g = \frac{1}{G}$, the above expression is further simplified as

$$\begin{aligned}\Sigma_{T_{ij}} - T_{ij} &= \frac{1}{G} \sum_{g=1}^G \left(\Sigma_{W_{ij}} - \frac{1}{n_g} \sum_{k \in I_g} (X_{ki} - \mu_{gi})(X_{kj} - \mu_{gj}) \right) \\ &\quad + \frac{1}{G} \sum_{g=1}^G (\mu_{gi}(\mu_{gj} - \bar{X}_{gj}) + \mu_{gj}(\mu_{gi} - \bar{X}_{gi})) + (\bar{X}_i \bar{X}_j - \mu_i \mu_j) \\ &= I_1 + I_2 + I_3\end{aligned}$$

For the final bound it remains to show that for each I_j there exist constants $c_{1j} > 0$ and $c_{2j} > 0$ such that

$$P(|I_j| \geq \epsilon) \leq c_{1j} \exp(-c_{2j} N \epsilon^2).$$

- $I_1 = \frac{1}{G} \sum_{g=1}^G \left(\Sigma_{W_{ij}} - \frac{1}{n_g} \sum_{k \in I_g} (X_{ki} - \mu_{gi})(X_{kj} - \mu_{gj}) \right).$

From Lemma A.3 in [Bickel & Levina \(2008\)](#) there exist constants $C_1 > 0$ and $C_2 > 0$ such that for $\epsilon < \epsilon_0$

$$P \left(\left| \frac{1}{n} \sum_{k=1}^n (Z_{ik} Z_{jk} - \sigma_{jk}) \right| \geq \epsilon \right) \leq C_1 \exp(-C_2 n \epsilon^2),$$

where Z_i are i.i.d $N(0, \Sigma)$ and σ_{ij} are elements of Σ . Let $\tilde{Z}_i = X_i - \mu_i$, where $\mu_i = \mu_g$ if observation i belongs to group g . By definition $\mathbb{E}(\tilde{Z}_i) = 0$ and \tilde{Z}_i are i.i.d $N(0, \Sigma_W)$. Note that I_1 can be rewritten as

$$I_1 = \frac{1}{N} \sum_{g=1}^G \sum_{k \in I_g} (\Sigma_{W_{ij}} - (X_{ki} - \mu_{gi})(X_{kj} - \mu_{gj})) = \frac{1}{N} \sum_{l=1}^N (\Sigma_{W_{ij}} - \tilde{Z}_{li} \tilde{Z}_{lj}).$$

Therefore

$$P(|I_1| \geq \epsilon) \leq P \left(\left| \frac{1}{N} \sum_{l=1}^N (\Sigma_{W_{ij}} - \tilde{Z}_{li} \tilde{Z}_{lj}) \right| \geq \epsilon \right) \leq C_1 \exp(-C_2 N \epsilon^2).$$

- $I_2 = \frac{1}{G} \sum_{g=1}^G (\mu_{gi}(\mu_{gj} - \bar{X}_{gj}) + \mu_{gj}(\mu_{gi} - \bar{X}_{gi})).$ It follows that

$$|I_2| \leq 2 \max \left(\left| \frac{1}{G} \sum_{g=1}^G \mu_{gi}(\mu_{gj} - \bar{X}_{gj}) \right|, \left| \frac{1}{G} \sum_{g=1}^G \mu_{gj}(\mu_{gi} - \bar{X}_{gi}) \right| \right).$$

Since the groups are independent,

$$\frac{1}{G} \sum_{g=1}^G \mu_{gi}(\mu_{gj} - \bar{X}_{gj}) \sim N \left(0, \frac{\sum_{g=1}^G \mu_{gi}^2}{NG} \sigma_j^2 \right).$$

Therefore,

$$P \left(\left| \frac{1}{G} \sum_{g=1}^G \mu_{gi}(\mu_{gj} - \bar{X}_{gj}) \right| \geq \epsilon \right) \leq 2 \exp(-c N \epsilon^2),$$

hence

$$P(|I_2| \geq \epsilon) \leq 4 \exp\left(-\frac{c}{4}N\epsilon^2\right).$$

- $I_3 = \bar{X}_i \bar{X}_j - \mu_i \mu_j$

Note that

$$\begin{aligned} \bar{X}_i \bar{X}_j - \mu_i \mu_j &= \sum_{g=1}^G \frac{1}{G} \bar{X}_{gi} \sum_{l=1}^G \frac{1}{G} \bar{X}_{lj} - \sum_{g=1}^G \frac{1}{G} \mu_{gi} \sum_{l=1}^G \frac{1}{G} \mu_{lj} \\ &= \sum_{g=1}^G \frac{1}{G} (\bar{X}_{gi} - \mu_{gi}) \sum_{l=1}^G \frac{1}{G} (\bar{X}_{lj} - \mu_{lj}) \\ &\quad + \mu_i \sum_{l=1}^G \frac{1}{G} (\bar{X}_{lj} - \mu_{lj}) + \mu_j \sum_{g=1}^G \frac{1}{G} (\bar{X}_{gi} - \mu_{gi}). \end{aligned}$$

Therefore

$$|I_3| \leq \max_{t \in \{i,j\}} \left| \frac{1}{G} \sum_{g=1}^G (\bar{X}_{gt} - \mu_{gt}) \right|^2 + 2 \max_g |\mu_g| \max_{t \in \{i,j\}} \left| \frac{1}{G} \sum_{g=1}^G (\bar{X}_{gt} - \mu_{gt}) \right|.$$

Since the groups are independent,

$$\frac{1}{G} \sum_{g=1}^G \bar{X}_{gj} \sim N\left(\frac{1}{G} \sum_{g=1}^G \mu_{gj}, \frac{\sigma_j^2}{nG}\right).$$

Therefore,

$$P\left(\max_{t \in \{i,j\}} \left| \frac{1}{G} \sum_{g=1}^G (\bar{X}_{gt} - \mu_{gt}) \right| \geq \epsilon\right) \leq 2 \cdot 2 \exp(-c_1 N \epsilon^2).$$

Note that if $\max_{t \in \{i,j\}} \left| \frac{1}{G} \sum_{g=1}^G (\bar{X}_{gt} - \mu_{gt}) \right| \leq k\epsilon$, then $|I_3| \leq k^2 \epsilon^2 + 2k\epsilon \max_g |\mu_g|$.

Choosing $k \leq \frac{1}{\max(\sqrt{2}, 2 \max_g \|\mu_g\|)}$ leads to $|I_3| \leq \frac{\epsilon^2 + \epsilon}{2} \leq \epsilon$ for small values of ϵ . Hence,

$$P(|I_3| \geq \epsilon) \leq 4 \exp(-cN\epsilon^2).$$

Combining I_1 - I_3 leads to $c_1 = C_1 + 4 + 4$. The results of lemma follow from the definition of $\|\cdot\|_\infty$ and the union bound. \square

Lemma 4. *Let $F_T = T_{A^c A}(T_{AA})^{-1} - \Sigma_{A^c A}(\Sigma_{AA})^{-1}$. Then there exists constant $c_3 > 0$ such that*

$$P(\|F_T\|_\infty \geq \epsilon \phi(\kappa + 1)(1 - \phi\epsilon)^{-1}) \leq 6(G + 1)ps \exp(-c_3 N s^{-2} \epsilon^2).$$

Proof. This result is a multi-group version of Lemma A2 in [Mai et al. \(2012\)](#). The proof follows the proof of Lemma A2 in [Mai et al. \(2012\)](#) and uses the results of Lemma 3. \square

Proof of Theorem 1

Proof. The proof follows the proof of Theorem 1 in [Mai et al. \(2012\)](#). For simplicity of illustration, the proof is given for the case $\pi_i = \frac{1}{G}$ and $n_i = \frac{N}{G}$. Extension to the general π_i and n_i is discussed in Appendix B.

Part 1. First we derive the conditions under which $\hat{V}(\lambda)_{A^c} = 0$. Let $T = W + DD^t$. From KKT conditions on $V = \hat{V}(\lambda)$, $TV - D + \lambda u = 0$, where u is the subgradient of $\sum_{i=1}^p \|v_i\|_2$ such that

$$u_j = \begin{cases} \frac{v_j}{\|v_j\|_2}, & \text{if } v_j \neq 0; \\ \in \{u : \|u\|_2 \leq 1\}, & \text{if } v_j = 0. \end{cases}$$

Partition T as

$$T = \begin{pmatrix} T_{AA} & T_{AA^c} \\ T_{A^cA} & T_{A^cA^c} \end{pmatrix}.$$

Then

$$\begin{aligned} T_{AA}V_A + T_{AA^c}V_{A^c} - D_A + \lambda u_A &= 0; \\ T_{A^cA}V_A + T_{A^cA^c}V_{A^c} - D_{A^c} + \lambda u_{A^c} &= 0. \end{aligned}$$

It follows that if $\|T_{A^cA}V_A - D_{A^c}\|_{\infty,2} \leq \lambda$, then $\hat{V}(\lambda)_{A^c} = 0$.

Next we derive the bounds for $\|T_{A^cA}V_A - D_{A^c}\|_{\infty,2}$. From KKT conditions

$$V_A = (T_{AA})^{-1}(D_A - \lambda u_A)$$

and therefore

$$\begin{aligned} T_{A^cA}V_A - D_{A^c} &= T_{A^cA}(T_{AA})^{-1}(D_A - \lambda u_A) - D_{A^c} \\ &= T_{A^cA}(T_{AA})^{-1}(D_A - \Delta_A) + T_{A^cA}(T_{AA})^{-1}\Delta_A - T_{A^cA}(T_{AA})^{-1}\lambda u_A - D_{A^c} + \Delta_{A^c} - \Delta_{A^c} \\ &= T_{A^cA}(T_{AA})^{-1}(D_A - \Delta_A) + T_{A^cA}(T_{AA})^{-1}\Delta_A - T_{A^cA}(T_{AA})^{-1}\lambda u_A + (\Delta_{A^c} - D_{A^c}) \\ &\quad - \Sigma_{A^cA}(\Sigma_{AA})^{-1}\Delta_A \\ &= (T_{A^cA}(T_{AA})^{-1} - \Sigma_{A^cA}(\Sigma_{AA})^{-1})\Delta_A + (\Delta_{A^c} - D_{A^c}) \\ &\quad + (T_{A^cA}(T_{AA})^{-1} - \Sigma_{A^cA}(\Sigma_{AA})^{-1})(D_A - \Delta_A) + \Sigma_{A^cA}(\Sigma_{AA})^{-1}(D_A - \Delta_A) \\ &\quad - (T_{A^cA}(T_{AA})^{-1} - \Sigma_{A^cA}(\Sigma_{AA})^{-1})\lambda u_A - \Sigma_{A^cA}(\Sigma_{AA})^{-1}\lambda u_A. \end{aligned}$$

Here we used the fact that $\tilde{\Psi} = \Sigma^{-1}\Delta P$ and $\tilde{\Psi} = \Sigma^{-1}\Delta$ have the same support. Also, since $\Sigma\tilde{\Psi} = \Delta$, it follows that $\Sigma_{AA}\tilde{\Psi}_A = \Delta_A$ and $\Sigma_{A^cA}\tilde{\Psi}_A = \Sigma_{A^cA}(\Sigma_{AA})^{-1}\Delta_A = \Delta_{A^c}$.

Denote $K = \Sigma_{A^cA}(\Sigma_{AA})^{-1}$, $F_T = T_{A^cA}(T_{AA})^{-1} - \Sigma_{A^cA}(\Sigma_{AA})^{-1}$, $F_{A^c} = \Delta_{A^c} - D_{A^c}$ and $F_A = \Delta_A - D_A$. Then

$$\begin{aligned} \|T_{A^cA}V_A - D_{A^c}\|_{\infty,2} &= \|F_T\Delta_A + F_{A^c} - F_TF_A - KF_A - F_T\lambda u_A - K\lambda u_A\|_{\infty,2} \\ &\leq \|F_T\|_{\infty,2}(\delta + \lambda) + \|F_{A^c}\|_{\infty,2} + \|F_T\|_{\infty}\|F_A\|_{\infty,2} + \|F_A\|_{\infty,2}\kappa + \lambda\kappa \\ &\leq \|F_T\|_{\infty}(\delta + \lambda) + \|F\|_{\infty,2}(1 + \|F_T\|_{\infty} + \kappa) + \lambda\kappa, \end{aligned}$$

where we used Lemma 1 and $\|F\|_{\infty,2} = \max(\|F_A\|_{\infty,2}, \|F_{A^c}\|_{\infty,2})$.

If $\|F_T\|_{\infty} \leq (\kappa + 1)\epsilon\phi(1 - \phi\epsilon)^{-1}$ and $\|F\|_{\infty,2} \leq \epsilon$ for some $\epsilon < \frac{1}{\phi}$ then

$$\|T_{A^cA}V_A - D_{A^c}\|_{\infty,2} \leq (\kappa + 1)\epsilon\phi(1 - \phi\epsilon)^{-1}(\delta + \lambda) + \epsilon(1 + (\kappa + 1)\epsilon\phi(1 - \phi\epsilon)^{-1} + \kappa) + \lambda\kappa.$$

Therefore $\|T_{A^c A} V_A - D_{A^c}\|_{\infty, 2} \leq \lambda$ if

$$(\kappa + 1)\epsilon\phi(1 - \phi\epsilon)^{-1}(\delta + \lambda) + \epsilon(1 + (\kappa + 1)\epsilon\phi(1 - \phi\epsilon)^{-1} + \kappa) + \lambda\kappa \leq \lambda,$$

which is equivalent to

$$\epsilon \leq \frac{\lambda(1 - \kappa)}{(\kappa + 1)(\phi\delta + 1) + 2\phi\lambda}.$$

Applying Lemma 2 and Lemma 3 leads to

$$P(\|T_{A^c A} V_A - D_{A^c}\|_{\infty} \leq \lambda) \geq 1 - c_1 p s \exp(-c_2 N s^{-2} \epsilon^2) - 2p(G - 1) \exp(-c_3 N \epsilon^2).$$

Part 2. From KKT conditions

$$\begin{aligned} \hat{V}(\lambda)_A &= T_{AA}^{-1}(D_A - \lambda u_A) = \Sigma_{AA}^{-1} \Delta + (T_{AA}^{-1} - \Sigma_{AA}^{-1} + \Sigma_{AA}^{-1})(D - \Delta) \\ &\quad + (T_{AA}^{-1} - \Sigma_{AA}^{-1})\Delta - \lambda(T_{AA}^{-1} - \Sigma_{AA}^{-1})u_A + \lambda \Sigma_{AA}^{-1} u_A. \end{aligned}$$

Denote $\nu_1 = \|T_{AA}^{-1} - \Sigma_{AA}^{-1}\|_{\infty}$. Then for any $j \in A$

$$\|\hat{v}_j\|_2 \geq \min_{i \in A} \|\tilde{\Psi}_i\|_2 - (\nu_1 + \phi)(\lambda + \|F_A\|_{\infty, 2}) - \nu_1 \delta.$$

Let $\nu_2 = \|T_{AA} - \Sigma_{AA}\|_{\infty}$. From the proof of Lemma A.2 in [Mai et al. \(2012\)](#):

$$\nu_1 < \phi^2 \nu_2 (1 - \phi \nu_2)^{-1}.$$

If $\nu_2 \leq \epsilon$ and $\|F_A\|_{\infty, 2} \leq \epsilon$, then $\|\hat{v}_j\|_2 > 0$ if

$$\epsilon < \frac{\tilde{\Psi}_{\min} - \lambda\phi}{\phi(1 + \phi\delta + \tilde{\Psi}_{\min})}.$$

It follows that

$$P(\|v_j\|_2 > 0 \text{ for all } j \in A) \geq 1 - c_1 s^2 \exp(-c_2 N s^{-2} \epsilon^2) - 2s(G - 1) \exp(-c_3 N \epsilon^2).$$

Part 3. From above,

$$\|\hat{V}(\lambda)_A - \tilde{V}\|_{\infty, 2} \leq (\nu_1 + \phi)(\|F_A\|_{\infty, 2} + \lambda) + \nu_1 \delta \leq (1 - \phi \nu_2)^{-1} \phi (\|F_A\|_{\infty, 2} + \lambda + \phi \nu_2 \delta).$$

If $\nu_2 \leq \epsilon$ and $\|F_A\|_{\infty, 2} \leq \epsilon$, then $\|\hat{V}(\lambda)_A - \tilde{V}\|_{\infty, 2} \leq 2\phi\lambda$ if

$$\epsilon \leq \frac{\lambda}{1 + \phi\delta + 2\phi\lambda}.$$

□

Appendix B

Extension to general $\pi_i > 0$ and n_i

Given that $n_i \sim \text{Bin}(N, \pi_i)$ for all $i = 1, \dots, G$, by Hoeffding inequality

$$P\left(\left|\pi_i - \frac{n_i}{N}\right| \geq \epsilon\right) \leq 2\exp(-2N\epsilon^2).$$

Hence

$$P\left(\left|\pi_i - \frac{n_i}{N}\right| \geq \frac{\pi_i}{2}\right) \leq 2\exp\left(-\frac{N\pi_i^2}{2}\right).$$

Consider event

$$\mathbf{A}_N = \bigcap_{i=1}^G \left\{ \frac{\pi_i N}{2} \leq n_i \leq \frac{3\pi_i N}{2} \right\}.$$

From above it follows that

$$P(\mathbf{A}_N) \geq 1 - 2G \exp\left(-\frac{N \min_i \pi_i^2}{2}\right).$$

The results of auxillary lemmas can be obtained by conditioning on Y and considering the probabilities on the event \mathbf{A}_N and its compliment. Since $P(\mathbf{A}_N)$ doesn't depend on p , s or ϵ , this conditioning does not affect the final rates.

Extension to sub-Gaussian sase

The normality assumption is employed to use the following concentration inequalities:

1. For any independent $X_1, \dots, X_N \sim N(\mu, \sigma^2)$

$$P(|\bar{X} - \mu| \geq \epsilon) \leq 2\exp(-cN\epsilon^2).$$

2. Lemma A.3 in [Bickel & Levina \(2008\)](#). For any independent $X_1, \dots, X_N \sim N_p(\mu, \Sigma)$ and $\epsilon \leq \epsilon_0$

$$P\left(\left|\frac{1}{N} \sum_{k=1}^N (X_{ki} - \mu_i)(X_{kj} - \mu_j) - \Sigma_{ij}\right| \geq \epsilon\right) \leq C_1 \exp(-C_2 N \epsilon^2).$$

Note that the first inequality remains true with the relaxation X_1, \dots, X_N are sub-Gaussian (see for example [Vershynin \(2010\)](#) and references therein). The second inequality remains true for sub-Gaussian random vectors X_1, \dots, X_N as long as

$$P\left(\left|\frac{1}{N} X^t Y - \mathbb{E}\left(\frac{1}{N} X^t Y\right)\right| > \epsilon\right) \leq C_1 \exp(-C_2 N \epsilon^2) \quad (14)$$

holds for zero-mean sub-Gaussian X and Y with independent individual components and some small $\epsilon < \epsilon_0$. If $X = Y$, then (14) is a special case of the Hanson-Wright inequality ([Rudelson & Vershynin, 2013](#)). Therefore the generalization of Hanson-Wright inequality for the case $X \neq Y$ is needed. If such a generalization can be obtained, then the results of Theorem 1 remain true with the relaxation $(X_i | Y_i = g)$ are sub-Gaussian.

References

- BICKEL, P. J. & LEVINA, E. (2004). Some theory for Fisher’s linear discriminant function, ‘naive Bayes’, and some alternatives when there are many more variables than observations. *Bernoulli* **10**, 989–1010.
- BICKEL, P. J. & LEVINA, E. (2008). Regularized estimation of large covariance matrices. *Annals of Statistics* **36**, 199–227.
- BOYD, S. P. & VANDENBERGHE, L. (2004). *Convex Optimization*. Cambridge: Cambridge Univ Press.
- CAI, T. & LIU, W. (2011). A direct estimation approach to sparse linear discriminant analysis. *Journal of the American Statistical Association* **106**, 1566–1577.
- CAI, T., LIU, W. & LUO, X. (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association* **106**, 594–607.
- CHAKRABORTY, S., GRUBER, T., BARRY, C. E., BOSHOFF, H. I. & RHEE, K. Y. (2013). Para-aminosalicylic acid acts as an alternative substrate of folate metabolism in *Mycobacterium tuberculosis*. *Science* **339**, 88–91.
- CLEMMENSEN, L., HASTIE, T., WITTEN, D. & ERSBØLL, B. (2011). Sparse discriminant analysis. *Technometrics* **53**, 406–413.
- DE CARVALHO, L. P. S., DARBY, C. M., RHEE, K. Y. & NATHAN, C. (2011). Nitazoxanide disrupts membrane potential and intrabacterial pH homeostasis of *Mycobacterium tuberculosis*. *ACS Medicinal Chemistry Letters* **2**, 849–854.
- DE CARVALHO, L. P. S., FISCHER, S. M., MARRERO, J., NATHAN, C., EHRT, S. & RHEE, K. Y. (2010). Metabolomics of *Mycobacterium tuberculosis* reveals compartmentalized co-catabolism of carbon substrates. *Chemistry & Biology* **17**, 1122–1131.
- DUDOIT, S., FRIDLYAND, J. & SPEED, T. P. (2002). Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American Statistical Association* **97**, 77–87.
- FAN, J. & FAN, Y. (2008). High dimensional classification using features annealed independence rules. *Annals of Statistics* **36**, 2605–2637.
- FAN, J., FENG, Y. & TONG, X. (2012). A road to classification in high dimensional space: the regularized optimal affine discriminant. *Journal of the Royal Statistical Society Ser. B* **74**, 745–771.
- FRIEDMAN, J. H. (1989). Regularized discriminant analysis. *Journal of the American Statistical Association* **84**, 165–175.
- GUO, Y., HASTIE, T. & TIBSHIRANI, R. (2007). Regularized linear discriminant analysis and its application in microarrays. *Biostatistics* **8**, 86–100.

- HASTIE, T., TIBSHIRANI, R. & FRIEDMAN, J. H. (2009). *The Elements of Statistical Learning, 2nd ed.* Springer Series in Statistics. New York: Springer.
- KOLAR, M. & LIU, H. (2013). Optimal feature selection in high-dimensional discriminant analysis. *ArXiv e-prints* **stat.ML**, arXiv:1306.6557.
- KRZANOWSKI, W. J., JONATHAN, P., MCCARTHY, W. V. & THOMAS, M. R. (1995). Discriminant analysis with singular covariance matrices: methods and applications to spectroscopic data. *Journal of the Royal Statistical Society Ser. C* **44**, 101–115.
- MAI, Q. & ZOU, H. (2013). A note on the connection and equivalence of three sparse linear discriminant analysis methods. *Technometrics* **55**, 243–246.
- MAI, Q., ZOU, H. & YUAN, M. (2012). A direct approach to sparse discriminant analysis in ultra-high dimensions. *Biometrika* **99**, 29–42.
- MARDIA, K. V., KENT, J. T. & BIBBY, J. M. (1979). *Multivariate Analysis*. New York: Academic Press.
- OBOZINSKI, G., WAINWRIGHT, M. J. & JORDAN, M. I. (2011). Support union recovery in high-dimensional multivariate regression. *Annals of Statistics* **39**, 1–47.
- PETHE, K., SEQUEIRA, P. C., AGARWALLA, S., RHEE, K., KUHEN, K., PHONG, W. Y., PATEL, V., BEER, D., WALKER, J. R., DURAISWAMY, J., JIRICEK, J., KELLER, T. H., CHATTERJEE, A., TAN, M. P., UJJINI, M., RAO, S. P. S., CAMACHO, L., BIFANI, P., MAK, P. A., MA, I., BARNES, S. W., CHEN, Z., PLOUFFE, D., THAYALAN, P., NG, S. H., AU, M., LEE, B. H., TAN, B. H., RAVINDRAN, S., NANJUNDAPPA, M., LIN, X., GOH, A., LAKSHMINARAYANA, S. B., SHOEN, C., CYNAMON, M., KREISWIRTH, B., DARTOIS, V., PETERS, E. C., GLYNNE, R., BRENNER, S. & DICK, T. (2010). A chemical genetic screen in Mycobacterium tuberculosis identifies carbon-source-dependent growth inhibitors devoid of in vivo efficacy. *Nature Communications* **1**, 57.
- RAMASWAMY, S., TAMAYO, P., RIFKIN, R., MUKHERJEE, S., YEANG, C. H., ANGELO, M., LADD, C., REICH, M., LATULIPPE, E., MESIROV, J. P., POGGIO, T., GERALD, W., LODA, M., LANDER, E. S. & GOLUB, T. R. (2001). Multiclass cancer diagnosis using tumor gene expression signatures. *Proceedings of the National Academy of Sciences of the United States of America* **98**, 15149–15154.
- RUDELSON, M. & VERSHYNIN, R. (2013). Hanson-Wright inequality and sub-gaussian concentration. *ArXiv e-prints* **math.PR**, arXiv:1306.2872.
- RUKHIN, A. L. (1992). Generalized Bayes estimators of a normal discriminant function. *Journal of Multivariate Analysis* **41**, 154–162.
- SEARLE, S. R. (2006). *Linear Models for Unbalanced Data*. New York: Wiley-Interscience.
- SHAO, J., WANG, Y., DENG, X. & WANG, S. (2011). Sparse linear discriminant analysis by thresholding for high dimensional data. *Annals of Statistics* **39**, 1241–1265.

- SIMON, N., FRIEDMAN, J. H., HASTIE, T. & TIBSHIRANI, R. (2013). A sparse-group Lasso. *Journal of Computational and Graphical Statistics* **22**, 231–245.
- TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B. & CHU, G. (2002). Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences of the United States of America* **99**, 6567–6572.
- TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B. & CHU, G. (2003). Class prediction by nearest shrunken centroids, with applications to DNA microarrays. *Statistical Science* **18**, 104–117.
- VERSHYNIN, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *ArXiv e-prints math.PR*, arXiv:1011.3027.
- WITTEN, D. M. & TIBSHIRANI, R. (2011). Penalized classification using Fisher’s linear discriminant. *Journal of the Royal Statistical Society Ser. B* **73**, 753–772.
- WU, M. C., ZHANG, L., WANG, Z., CHRISTIANI, D. C. & LIN, X. (2009). Sparse linear discriminant analysis for simultaneous testing for the significance of a gene set/pathway and gene selection. *Bioinformatics* **25**, 1145–1151.
- XU, P., BROCK, G. N. & PARRISH, R. S. (2009). Modified linear discriminant analysis approaches for classification of high-dimensional microarray data. *Computational Statistics & Data Analysis* **53**, 1674–1687.
- YUAN, M. & LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Ser. B* **68**, 49–67.